

CLAIMS

We claim:

1. A method for data distribution, comprising:
 - distributing logical addresses among an initial set of storage devices so as to provide a balanced access to the devices;
 - transferring the data to the storage devices in accordance with the logical addresses;
 - adding an additional storage device to the initial set, thus forming an extended set of the storage devices comprising the initial set and the additional storage device; and
 - redistributing the logical addresses among the storage devices in the extended set so as to cause a portion of the logical addresses to be transferred from the storage devices in the initial set to the additional storage device, while maintaining the balanced access and without requiring a substantial transfer of the logical addresses among the storage devices in the initial set.
2. A method according to claim 1, wherein redistributing the logical addresses comprises no transfer of the logical addresses between the storage devices in the initial set.
3. A method according to claim 1, wherein distributing the logical addresses comprises applying a consistent hashing function to the initial set of storage devices so as to determine respective initial locations of the logical addresses among the initial set, and wherein redistributing the logical addresses comprises applying the consistent hashing function to the extended set of storage devices so as to determine respective subsequent locations of the logical addresses among the extended set.

4. A method according to claim 1, wherein distributing the logical addresses comprises applying a randomizing function to the initial set of storage devices so as to determine respective initial locations of the logical
5 addresses among the initial set, and wherein redistributing the logical addresses comprises applying the randomizing function to the extended set of storage devices so as to determine respective subsequent locations of the logical addresses among the extended
10 set.

5. A method according to claim 1, wherein at least one of the storage devices comprises a fast access time memory.

6. A method according to claim 1, wherein at least one
15 of the storage devices comprises a slow access time mass storage device.

7. A method according to claim 1, wherein the storage devices have substantially equal capacities, and wherein distributing the logical addresses comprises distributing
20 the logical addresses substantially evenly among the initial set, and wherein redistributing the logical addresses comprises redistributing the logical addresses substantially evenly among the extended set.

8. A method according to claim 1, wherein a first
25 storage device comprised in the storage devices has a first capacity different from a second capacity of a second storage device comprised in the storage devices, and wherein distributing the logical addresses comprises distributing the logical addresses substantially
30 according to a ratio of the first capacity to the second capacity, and wherein redistributing the logical addresses comprises redistributing the logical addresses substantially according to the ratio.

9. A method according to claim 1, wherein distributing the logical addresses comprises allocating a specific logical address to a first storage device and to a second storage device, the first and second storage devices comprising different storage devices, and wherein storing the data comprises storing a first copy of the data on the first storage device and a second copy of the data on the second storage device.

10. A method according to claim 1, and comprising writing the data from a host external to the storage devices, and reading the data to the external host from the storage devices.

11. A method for distributing data, comprising:

15 distributing logical addresses among an initial set of storage devices so as provide a balanced access to the devices;

transferring the data to the storage devices in accordance with the logical addresses;

20 removing a surplus device from the initial set, thus forming a depleted set of the storage devices comprising the initial storage devices less the surplus storage device; and

25 redistributing the logical addresses among the storage devices in the depleted set so as to cause logical addresses of the surplus device to be transferred to the depleted set, while maintaining the balanced access and without requiring a substantial transfer of logical addresses among the storage devices in the depleted set.

30 12. A method according to claim 11, wherein redistributing the logical addresses comprises no transfer of the logical addresses to the storage devices in the depleted set apart from the logical addresses of

the surplus device.

13. A method according to claim 11, wherein distributing the logical addresses comprises applying a consistent hashing function to the initial set of storage devices so 5 as to determine respective initial locations of the logical addresses among the initial set, and wherein redistributing the logical addresses comprises applying the consistent hashing function to the depleted set of storage devices so as to determine respective subsequent 10 locations of the logical addresses among the depleted set.

14. A method according to claim 11, wherein distributing the logical addresses comprises applying a randomizing function to the initial set of storage devices so as to 15 determine respective initial locations of the logical addresses among the initial set, and wherein redistributing the logical addresses comprises applying the randomizing function to the depleted set of storage devices so as to determine respective subsequent 20 locations of the logical addresses among the depleted set.

15. A method according to claim 11, wherein at least one of the storage devices comprises a fast access time memory.

25 16. A method according to claim 11, wherein at least one of the storage devices comprises a slow access time mass storage device.

17. A method according to claim 11, wherein the storage devices have substantially equal capacities, and wherein 30 distributing the logical addresses comprises distributing the logical addresses substantially evenly among the initial set, and wherein redistributing the logical addresses comprises redistributing the logical addresses

substantially evenly among the depleted set.

18. A method according to claim 11, wherein a first storage device comprised in the storage devices has a first capacity different from a second capacity of a 5 second storage device comprised in the storage devices, and wherein distributing the logical addresses comprises distributing the logical addresses substantially according to a ratio of the first capacity to the second capacity, and wherein redistributing the logical 10 addresses comprises redistributing the logical addresses substantially according to the ratio.

19. A method according to claim 11, wherein distributing the logical addresses comprises allocating a specific logical address to a first storage device and to a second 15 storage device, the first and second storage devices comprising different storage devices, and wherein storing the data comprises storing a first copy of the data on the first storage device and a second copy of the data on the second storage device.

20. A method according to claim 11, and comprising writing the data from a host external to the storage devices, and reading the data to the external host from the storage devices.

21. A method for distributing data among a set of 25 storage devices, comprising:

applying a consistent hashing function to the set so as to allocate logical addresses to respective primary storage devices of the set and so as to provide a balanced access to the devices;

30 forming subsets of the storage devices by subtracting the respective primary storage devices from the set;

applying the consistent hashing function to the

subsets so as to allocate the logical addresses to respective secondary storage devices of the subsets while maintaining the balanced access to the devices; and

5 storing the data on the respective primary storage devices and a copy of the data on the respective secondary storage devices in accordance with the logical addresses.

22. A method for distributing data among a set of storage devices, comprising:

10 applying a randomizing function to the set so as to allocate logical addresses to respective primary storage devices of the set and so as to provide a balanced access to the devices;

15 forming subsets of the storage devices by subtracting the respective primary storage devices from the set;

20 applying the randomizing function to the subsets so as to allocate the logical addresses to respective secondary storage devices of the subsets while maintaining the balanced access to the devices; and

storing the data on the respective primary storage devices and a copy of the data on the respective secondary storage devices in accordance with the logical addresses.

25 23. A data distribution system, comprising:

an initial set of storage devices among which are distributed logical addresses so as provide a balanced access to the devices, and wherein data is stored in accordance with the logical addresses; and

30 an additional storage device to the initial set, thus forming an extended set of the storage devices comprising the initial set and the additional storage device, the logical addresses being redistributed among the storage devices in the extended set so as to cause a

portion of the logical addresses to be transferred from the storage devices in the initial set to the additional storage device, while maintaining the balanced access and without requiring a substantial transfer of the logical
5 addresses among the storage devices in the initial set.

24. A system according to claim 23, and wherein the logical addresses are redistributed so that there is no transfer of the logical addresses between the storage devices in the initial set.

10 25. A system according to claim 23, wherein the distributed logical addresses are determined by applying a consistent hashing function to the initial set of storage devices so as to determine respective initial locations of the logical addresses among the initial set,
15 and wherein redistributing the logical addresses comprises applying the consistent hashing function to the extended set of storage devices so as to determine respective subsequent locations of the logical addresses among the extended set.

20 26. A system according to claim 23, wherein the distributed logical addresses are determined by applying a randomizing function to the initial set of storage devices so as to determine respective initial locations of the logical addresses among the initial set, and
25 wherein redistributing the logical addresses comprises applying the randomizing function to the extended set of storage devices so as to determine respective subsequent locations of the logical addresses among the extended set.

30 27. A system according to claim 23, wherein at least one of the storage devices comprises a fast access time memory.

28. A system according to claim 23, wherein at least one

of the storage devices comprises a slow access time mass storage device.

29. A system according to claim 23, wherein the storage devices have substantially equal capacities, and wherein
5 the distributed logical addresses are distributed substantially evenly among the initial set, and wherein redistributing the logical addresses comprises redistributing the logical addresses substantially evenly among the extended set.

10 30. A system according to claim 23, wherein a first storage device comprised in the storage devices has a first capacity different from a second capacity of a second storage device comprised in the storage devices, and wherein the distributed logical addresses are
15 distributed substantially according to a ratio of the first capacity to the second capacity, and wherein redistributing the logical addresses comprises redistributing the logical addresses substantially according to the ratio.

20 31. A system according to claim 23, wherein the distributed logical addresses comprise a specific logical address allocated to a first storage device and to a second storage device, the first and second storage devices comprising different storage devices, and wherein
25 storing the data comprises storing a first copy of the data on the first storage device and a second copy of the data on the second storage device.

32. A system according to claim 23, and comprising a memory having a table wherein is stored a correspondence
30 between a plurality of the logical addresses and a specific storage device in the initial set, wherein the plurality of the logical addresses are related to each other by a mathematical relation.

33. A data distribution system, comprising:

an initial set of storage devices among which are distributed logical addresses so as provide a balanced access to the devices, and wherein data is stored in accordance with the logical addresses; and

a depleted set of storage devices, formed by subtracting a surplus storage device from the initial set, the logical addresses being redistributed among the storage devices in the depleted set so as to cause logical addresses of the surplus device to be transferred to the depleted set, while maintaining the balanced access and without requiring a substantial transfer of the logical addresses among the storage devices in the depleted set.

34. A system according to claim 33; wherein redistributing the logical addresses comprises no transfer of the logical addresses to the storage devices in the depleted set apart from the logical addresses of the surplus device.

35. A system according to claim 33, wherein the distributed logical addresses are determined by applying a consistent hashing function to the initial set of storage devices so as to determine respective initial locations of the logical addresses among the initial set, and wherein redistributing the logical addresses comprises applying the consistent hashing function to the depleted set of storage devices so as to determine respective subsequent locations of the logical addresses among the depleted set.

36. A system according to claim 33, wherein the distributed logical addresses are determined by applying a randomizing function to the initial set of storage devices so as to determine respective initial locations

of the logical addresses among the initial set, and wherein redistributing the logical addresses comprises applying the randomizing function to the depleted set of storage devices so as to determine respective subsequent 5 locations of the logical addresses among the depleted set.

37. A system according to claim 33, wherein at least one of the storage devices comprises a fast access time memory.

10 38. A system according to claim 33, wherein at least one of the storage devices comprises a slow access time mass storage device.

39. A system according to claim 33, wherein the storage devices have substantially equal capacities, and wherein 15 the distributed logical addresses are distributed substantially evenly among the initial set, and wherein redistributing the logical addresses comprises redistributing the logical addresses substantially evenly among the depleted set.

20 40. A system according to claim 33, wherein a first storage device comprised in the storage devices has a first capacity different from a second capacity of a second storage device comprised in the storage devices, and wherein the distributed logical addresses are 25 distributed substantially according to a ratio of the first capacity to the second capacity, and wherein redistributing the logical addresses comprises redistributing the logical addresses substantially according to the ratio.

30 41. A system according to claim 33, wherein the distributed logical addresses comprise a specific logical address allocated to a first storage device and a second storage device, the first and second storage devices

comprising different storage devices, and wherein storing the data comprises storing a first copy of the data on the first storage device and a second copy of the data on the second storage device.

5 42. A system according to claim 33, and comprising a memory having a table wherein is stored a correspondence between a plurality of logical addresses and a specific storage device in the initial set, wherein the plurality of logical addresses are related to each other by a
10 mathematical relation.

43. A data distribution system, comprising:

a set of data storage devices to which is applied a consistent hashing function so as to allocate logical addresses to respective primary storage devices of the
15 set and so as to provide a balanced access to the devices; and

subsets of the storage devices formed by subtracting the respective primary storage devices from the set, the consistent hashing function being applied to the subsets
20 so as to allocate the logical addresses to respective secondary storage devices of the subsets while maintaining the balanced access to the devices, data being stored on the respective primary storage devices and a copy of the data being stored on the respective
25 secondary storage devices in accordance with the logical addresses.

44. A data distribution system, comprising:

a set of data storage devices to which is applied a randomizing function so as to allocate logical addresses
30 to respective primary storage devices of the set and so as to provide a balanced access to the devices; and

subsets of the storage devices formed by subtracting the respective primary storage devices from the set, the

randomizing function being applied to the subsets so as to allocate the logical addresses to respective secondary storage devices of the subsets while maintaining the balanced access to the devices, data being stored on the 5 respective primary storage devices and a copy of the data being stored on the respective secondary storage devices in accordance with the logical addresses.